

Claims

We claim:

1 SUB A27 1. A mapping table for referencing rows of a primary B+tree, the mapping table
2 comprising:
3 a row for each row of the primary B+tree.

2. The mapping table according to claim 1, wherein each row of the mapping table
comprises a primary key value from the primary B+tree.

3. The mapping table according to claim 1, wherein the mapping table provides
one-to-one mapping between primary keys of the primary B+tree structure and physical
row identifiers of the mapping table.

1 4. The mapping table according to claim 1, wherein each row of the mapping table
2 comprises a guess-DBA, database block address of a leaf block of the primary B+tree,
3 where the corresponding primary B+tree row is likely to be found.

1 / 5. A primary B+tree, comprising:
2 mapping table row identifiers stored in each row of the primary B+tree, the mapping table

3 row identifiers comprising a physical row identifier of a corresponding mapping table
4 row.

1 6. The primary B+tree according to claim 5, wherein the mapping table row
2 identifiers are stored at a fixed offset from a beginning of each row of the primary B+tree.

1 / 7. An auxiliary structure for a primary B+tree, the auxiliary structure comprising:
2 row identifiers of corresponding mapping table rows, the row identifiers referring to a
3 primary B+tree row.

1 / 8. A method for loading/populating a primary B+tree having an associated
2 mapping table, the method comprising:
3 generating a row of the mapping table for each row of the primary B+tree; and
4 storing in each row of the mapping table a row identifier for a corresponding row
5 of the primary B+tree, the row identifier comprising a primary key column value for each
6 row of the primary B+tree and a guess-DBA.

1 / 9. A method for maintaining a circular dependency between a mapping table row
2 and a primary B+tree row, the method comprising:
3 computing a length of a mapping table row based upon a length of a primary key
4 and an overhead of a guess-DBA;

5 utilizing the computed length to identify a mapping table block that can
6 accommodate the row;
7 reserving a slot in the identified mapping table block, wherein an address of the
8 block and a reserved slot form a mapping table physical row identifier;
9 inserting a primary B+tree row containing the physical row identifier into the
10 primary B+tree;
11 utilizing a leaf block address of the primary B+tree row to construct a row of the
12 mapping table; and
13 inserting the mapping table row in the reserved slot.

10. The method according to claim 9, further comprising:
carrying out a partition maintenance operation on the primary B+tree; and
rebuilding the mapping table after the partition maintenance.

11. The method according to claim 9, further comprising:
carrying out a partition maintenance operation on the primary B+tree; and
maintain the mapping table during the partition maintenance.

12. The method according to claim 9, further comprising:
carrying out a partition maintenance operation on the primary B+tree and
rebuilding the mapping table on-line.

13. A computer program product for performing a process for indexing a primary B+tree, the computer program product comprising:
a computer readable medium; and
computer program instructions, recorded on the computer readable medium, executable by a processor, for performing the steps of:
generating a row of a mapping table for each row of the primary B+tree; and
storing in each row of the mapping table a row identifier for a corresponding row of the primary B+tree, the identifier comprising a primary key column value and a guess-database address for each row of the primary B+tree.

14. A system for performing a process for indexing a primary B+tree, the system comprising:
a processor operable to execute computer program instructions; and
a memory operable to store computer program instructions executable by the processor, for performing the steps of:
generating a row of a mapping table for each row of the primary B+tree; and
storing in each row of the mapping table a row identifier for a corresponding row of the primary B+tree, the identifier comprising a primary key column value and a guess-database address for each row of the primary B+tree.